

2020

- [January 2020](#)
- [February 2020](#)
- [March 2020](#)
- [April 2020](#)
- [May 2020](#)
- [September 2020](#)

January 2020

January 4th

I received the dell precision rack I bought on ebay. The performance is quite good: on linpack the installed e5-2623v3's get 260ish Gflops. I think, moving forward, I would like to augment this with more memory and more cores, especially if this becomes a datascience box and I repurpose the z620 for something else.

I also learned that this machine has full support for nvme boot but requires pcie to nvme adapters in order to do so. Dell sold this with a quad nvme drive called the Dell Ultra-Speed Drive Quad NVMe

January 7th

I attempted to switch my e585 thinkpad from the nvme drive to a sata drive so that I could suspend again. However, this did not fix the suspend issues I'm having. Now, it occasionally gives me a white screen instead of a black screen but the end result is the same: it requires a hard reset using the power button. I removed the sata ssd and put the nvme back in because it's faster and also the only computer I have that can use the nvme drive at the moment. since both cause the same errors, I'm just sticking with what I had. I put the ssd in the precision rack computer, though it still has opensuse installed on it for now.

January 8th

Currently, the disk in the small optiplex server is just floating in there which is a not so great setup. I 3d printed a caddy to put the hard drive in but the way the cable lengths are set up, it

wouldn't be quite long enough to reach the power on both the ssd and hard drive at the same time.

I need to flip the holes in the stl file so that the hard drive is in upside down from how it normally is. This will give me the slack I require in the sata power cable to feed both the ssd and hard drive at the same time.

January 9th

iDRAC and proxmox are setup on the dell precision rack 7910 that I got on the 4th. It works quite well. I 3d printed caddies to include in there. One thing that is kind of strange is that there's a slightly rattly sound to one of the fans. Perhaps I'll have to find a replacement. I'm not sure which fan it is.

January 10th

Bookstack upgrade

Upgraded bookstack, apparently I hadn't upgraded it since the initial install in 2018. The upgrade appears to have gone without a hitch.

Container migration onto moja

I am unable to migrate containers from mbili and tatu onto moja because local-lvm does not exist as a storage device there. I think I will take the 1tb ssd in that node and make it a lvm vg to enable movement of things onto there.

Storage migration of jellyfin

Overnight, I ran a migration of jellyfin from lizardfs to zfs on tatu. lizardfs is great but because jellyfin only supports sqlite, which uses many small reads, I can't leverage my postgres server running on flash. The small reads for sqlite are the worst case scenario for lizardfs and as such the server runs very sluggish. After the migration from lizardfs to zfs, it's very snappy. However, this reduces my ability to migrate it in the future. lizardfs to it's credit has already balanced files away from tatu since lizardfs is sharing the zpool on that machine with jellyfin now.

hot backup

I should probably setup hot backups with a limit of 1 that run every night and backup to lizardfs, this would allow quicker migration despite my diminishing use of lizardfs for actually hosting containers.

I set this up on January 13th

it does backups every day except saturday since I have an existing backup job running every saturday. The shared storage created has a 1 backup limit so this should only hold a backup for the last day.

February 2020

<https://www.proxmox.com/en/training/video-tutorials/item/bond-configuration>

Infiniband notes

- Infiniband is a good idea for multi-node gpu training as it can use RDMA. This allows exchanges between the gpus on each node without first copying the data from the gpu into system ram.
- QDR:
 - 40 gbps per port
 - 1 micro second latency
 - \$40 for a dual port card
- Bonding of infiniband interfaces is not possible?

Purchases

- Bought 3 quad intel nics for \$35 on ebay. This should help with lizardfs bandwidth.
- Bought an m40 gpu with 24gb of vram.
 - Needs a fan and the 3d printed bracket for adding fans to the passive card.

March 2020

March 4 2020

Trying to get tesla m40 into z620

UEFI is required!! I converted the z620 machine (mwanafunzi) from legacy boot to uefi boot and switched the gpu from legacy to efi support. [Here](#) is the pastebin output from that.

This worked very well, with the tesla drivers, I got the m40 gpu to show up with all 24 GB of VRAM.

```
      .-/+00SSSS00+/- .
      `:+SSSSSSSSSSSSSSSSSS+:`
      -+SSSSSSSSSSSSSSSSSSySSSS+-
      .OSSSSSSSSSSSSSSSSSSdMMMNySSSSO.
      /SSSSSSSSSSShdmmNNmmyNMMMMhSSSSSS/
      +SSSSSSSSShmydMMMMMMMMNddddySSSSSSSS+
      /SSSSSSSSShNMMMyhhyyyhmNMMMNhSSSSSSSS/
      .SSSSSSSSdMMMNhSSSSSSSSShNMMMdSSSSSSSS.
+SSSShhhyNMMNySSSSSSSSSSSyNMMMySSSSSSSS+
ossyNMMMNyMMhSSSSSSSSSSShmmhSSSSSSSO
ossyNMMMNyMMhSSSSSSSSSSShmmhSSSSSSSO
+SSSShhhyNMMNySSSSSSSSSSSyNMMMySSSSSSSS+
.SSSSSSSdMMMNhSSSSSSSSShNMMMdSSSSSSSS.
/SSSSSSSShNMMMyhhyyyhdNMMMNhSSSSSSSS/
+SSSSSSSSdmydMMMMMMMMMdddySSSSSSSS+
/SSSSSSSSShdmNNNNmyNMMMMhSSSSSS/
.OSSSSSSSSSSSSSSSSSSdMMMNySSSSO.
  -+SSSSSSSSSSSSSSSSSSyyySSSS+-
    `:+SSSSSSSSSSSSSSSSSS+:`
      .-/+00SSSS00+/- .

root@mwanafunzi
-----
OS: Ubuntu 18.04.4 LTS x86_64
Host: HP Z620 Workstation
Kernel: 5.3.0-40-generic
Uptime: 25 secs
Packages: 3166
Shell: bash 4.4.20
Terminal: run-parts
CPU: Intel Xeon E5-2650 0 (16) @ 2.800GHz
GPU: NVIDIA Tesla M40
GPU: NVIDIA GeForce GTX 1070
Memory: 716MiB / 48123MiB
```

I tested this out with allennlp and it worked quite well. The m40 trained an rnn about as fast as my 1070's. This is the worst case for this comparison because the 1070 has a higher clock rate but fewer cuda cores and rnn's are difficult to parallelize.

The additional vram allowed me to go up to 128 for the batch size.

However, my cooling solution for the m40 was insufficient. After about 1 epoch of training (5 minutes of heavy usage) the temp exceeded 80 C and I had to end the workload. I was using a single NF 4x20A fan from noctua. However, these only provide 5 CFM of airflow. I purchased a 2 pack of delta 40mm fans that achieve 10 CFM each. This should be enough airflow for operation. While the noise level is going up, it is only increasing from 17 dba to 35 dba (per the documentation for the respective products). since the z620 case fans are about 35 dba, this noise difference shouldn't be very noticable.

AMD GPU build (RX580)

batch_size 32 tensorflow 1.14

resnet50: 49.93 images per second

resnet 152: 20.83 images per second

inception v3: 20.03 images per second

April 2020

I really want to increase the size of the sdd in the z620 so that I can snapshot any time I do an apt upgrade for easy rollbacks.

I think everything necessary to remove the drive from mwanafunzi has been done: the virtual group has been deleted along with the logical volumes inside it.

The issue remaining is how to move all the user accounts over. I've moved one user account over before but movign all of them may cause conflicts in UIDs. I think this can be done by first adding the user accounts and then copying the `/etc/passwd` and `/etc/shadow` entries over, changing the UIDs to match the newly created ones.

I need to make sure to take out the 500GB HDD when I do the reinstall so that the installation process doesn't identify the existing boot partition and keep it on that drive. I want to get everything onto a single disk this time to reduce the likelihood of failure (since I'm not using RAID for the root partition)

May 2020

The hard drive in mbili is showing some read errors. I need to 3d print a new drive caddy for that computer.

Turns out I can use the caddy I printed before and jam it in there. Doesn't seem to be throwing errors anymore.

September 2020

Trying to build pytorch with hip/rocm support on ubuntu 20.04 and rocm 3.7.

```
CMake Error: The following variables are used in this project, but they are set to NOTFOUND.
Please set them or make sure they are set and tested correctly in the CMake files:
GLOO_HIP_HCC_LIBRARIES    linked by target "gloo_hip" in directory
/home/kenneth/build2/pytorch_rocm/third_party/gloo/gloo
PYTORCH_HIP_HCC_LIBRARIES  linked by target "cl0_hip" in directory
/home/kenneth/build2/pytorch_rocm/cl0/hip    linked by target "caffe2_nvrtc" in directory
/home/kenneth/build2/pytorch_rocm/caffe2     linked by target "torch_hip" in directory
/home/kenneth/build2/pytorch_rocm/caffe2
ROCM_HIPRTC_LIB          linked by target "caffe2_nvrtc" in directory
/home/kenneth/build2/pytorch_rocm/caffe2     linked by target "torch_hip" in directory
/home/kenneth/build2/pytorch_rocm/caffe2
-- Configuring incomplete, errors occurred!See also
"/home/kenneth/build2/pytorch_rocm/build/CMakeFiles/CMakeOutput.log".See also
"/home/kenneth/build2/pytorch_rocm/build/CMakeFiles/CMakeError.log".Traceback (most recent
call last):
  File "setup.py", line 732, in <module>
    build_deps()
  File "setup.py", line 311, in build_deps
    build_caffe2(version=version, File
"/home/kenneth/build2/pytorch_rocm/tools/build_pytorch_libs.py", line 54, in build_caffe2
cmake.generate(version, File
"/home/kenneth/build2/pytorch_rocm/tools/setup_helpers/cmake.py", line 329, in generate
self.run(args, env=my_env) File
"/home/kenneth/build2/pytorch_rocm/tools/setup_helpers/cmake.py", line 141, in run
check_call(command, cwd=self.build_dir, env=env) File "/usr/lib/python3.8/subprocess.py",
line 364, in check_call
    raise CalledProcessError(retcode, cmd)
```